

Term Premia and Interest Rate Forecasts in Affine Models

Gregory R. Duffee
Haas School of Business
University of California—Berkeley
Visiting Scholar, Federal Reserve Bank of San Francisco

First version November 1998
This version November 2000
Previously circulated under the title
Forecasting Future Interest Rates: Are Affine Models Failures?

Abstract

I find that the standard class of affine models produces poor forecasts of future changes in Treasury yields. Better forecasts are generated by assuming that yields follow random walks. The failure of these models is driven by one of their key features: The compensation that investors receive for facing risk is a multiple of the variance of the risk. This means that risk compensation cannot vary independently of interest rate volatility. I also describe and empirically estimate a class of models that is broader than the standard affine class. These ‘essentially affine’ models retain the tractability of the usual models, but allow the compensation for interest rate risk to vary independently of interest rate volatility. This additional flexibility proves useful in forming accurate forecasts of future yields.

The most recent version of this paper is at <http://www.haas.berkeley.edu/~duffee>. Address correspondence to the University of California, Haas School of Business, 545 Student Services Building #1900, Berkeley, CA 94720. Phone: 510-642-1435. Email address: duffee@haas.berkeley.edu. I thank Mark Fisher for many helpful conversations and Jonathan Berk, Rob Bliss, Qiang Dai, Darrell Duffie, an anonymous referee, and seminar participants at Berkeley, Yale, Stanford, and an NBER Asset Pricing meeting for useful comments. The views expressed in this paper are the author’s and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

Can we use finance theory to tell us something about the empirical behavior of Treasury yields that we don't already know? In particular, can we sharpen our ability to predict future levels of yields? A long-established fact about Treasury yields is that the current term structure contains information about future term structures. For example, long-maturity bond yields tend to fall over time when the slope of the yield curve is steeper than usual. These predictive relations are based exclusively on the time-series behavior of yields. We know from finance theory that the cross-sectional and time-series behavior of the term structure must be linked in an internally consistent way in order to avoid arbitrage opportunities. In principle, imposing this restriction should allow us to exploit more of the information in the current term structure, and thus improve forecasts. But in practice, existing no-arbitrage models impose other restrictions for the sake of tractability, thus their value as forecasting tools is *a priori* unclear.

I examine the forecasting ability of the affine class of term structure models. By 'affine', I refer to models where zero-coupon bond yields, their physical (i.e., true) dynamics, and their equivalent martingale (i.e., risk-adjusted) dynamics are all affine functions of an underlying state vector. A variety of non-affine models have been developed, but the tractability and apparent richness of the affine class has led the finance profession to focus most of its attention on such models.

Although forecasting future yields is important in its own right, a model that is consistent with finance theory and produces accurate forecasts can make a deeper contribution to finance. It should allow us to address a key issue: explaining the well-documented time-variation in expected returns to assets. In the context of the term structure, explaining time-variation in expected returns means explaining the failure of the expectations hypothesis of interest rates. Put differently, we would like to have an intuitive explanation for the positive correlation between the slope of the yield curve and expected excess returns to long bonds. If a model produces poor forecasts of future yields (and thus poor forecasts of future bond prices), it is unlikely that the model can shed light on the economics underlying the failure of the expectations hypothesis.

The first main conclusion reached in this paper is that the class of affine models studied most extensively to date fails at forecasting. I refer to this class, which includes multifactor generalizations of both Vasicek (1977) and Cox, Ingersoll, and Ross (1985), and is extensively analyzed in Dai and Singleton (2000), as "completely affine." I fit general three-factor completely affine models to the Treasury term structure (with maturities ranging from three months to ten years) over the period 1952 through 1994. Yield forecasts produced with these estimated models are typically worse than forecasts produced by simply assuming yields follow random walks. This conclusion holds for both in-sample

forecasts and out-of-sample (1995 through 1998) forecasts.

Even more damning is the way in which the estimated models fail. They produce yield forecast errors that are strongly negatively correlated with the slope of the yield curve. In other words, the models fail to replicate the key empirical relation between expected returns and the slope of the yield curve; their underestimation of expected excess returns to long bonds tends to be largest when the slope of the term structure is steep.

This failure is a consequence of two features of the Treasury term structure, combined with a restriction built into completely affine models. The first feature is that Treasury yields vary widely over time around both sides of their (sample) means. Another way to say this is that we observe a wide variety of term structure shapes in the data. The second feature is that across the entire maturity spectrum, the unconditional mean excess return to bonds is small relative to the variation in conditional mean excess returns. While the average return to Treasury bonds is not much greater than zero, the slope of the term structure predicts a relatively large amount of variation in excess returns to bonds. One implication of this second feature is that, as noted by Fama and French (1993), the sign of predicted excess returns to Treasury bonds changes over time.

Completely affine models do not simultaneously reproduce these two features of term structure behavior. The key restriction in these models is that compensation for risk is a fixed multiple of the variance of the risk. This structure ensures that the models satisfy a requirement of no-arbitrage: Risk compensation goes to zero if risk goes to zero. But because variances are nonnegative, this structure also places an important limitation on the time-series behavior of the compensation that investors expect to receive for facing a given risk. The compensation is bounded by zero, therefore it cannot switch sign over time.

As will be made clear in the paper, the only way this framework can produce expected excess returns with low means and high volatilities is for some underlying factors driving the term structure to be highly positively skewed. But this strong positive skewness limits the flexibility of the model to fit a wide variety of term structure shapes. Thus completely affine models can fit either of these features of Treasury yields, but not both simultaneously.

All is not lost, however. The second main conclusion of this paper is that the completely affine class can be extended to break the link between risk compensation and interest rate volatility. This extension from the completely affine class to the “essentially affine” class described here is costless, in the sense that the affine time-series and cross-sectional properties of bond prices are preserved in essentially affine models. The existence of extensions to the completely affine class is not new (Chacko, 1997, constructs a general equilibrium example), but this paper is the first to describe and empirically investigate a

general, very tractable extension to completely affine term structure models. I find that essentially affine models can produce more accurate yield forecasts than completely affine models, both in-sample and out-of-sample. However, there is a tradeoff between flexibility in forecasting future yields and flexibility in fitting interest rate volatility.

The paper is organized as follows. The structure of affine models is discussed in detail in Section 1. Section 2 explains intuitively why completely affine models work poorly. Section 3 describes the estimation technique. Section 4 presents empirical results and Section 5 concludes.

1. Affine Models of the Term Structure

1.1. Affine bond pricing

The core of affine term structure models is the framework of Duffie and Kan (1996). Their model, which is summarized here, describes the evolution of bond prices under the equivalent martingale measure. Uncertainty is generated by n Brownian motions, $\tilde{W}_t \equiv (\tilde{W}_{t,1}, \dots, \tilde{W}_{t,n})'$. There are n state variables, denoted $X_t \equiv (X_{t,1}, \dots, X_{t,n})'$. The instantaneous nominal interest rate, denoted r_t , is affine in these state variables:

$$r_t = \delta_0 + \delta X_t,$$

where δ_0 is a scalar and δ is an n -vector. The evolution of the state variables under the equivalent martingale measure is

$$dX_t = [(K\theta)^Q - K^Q X_t]dt + \Sigma S_t d\tilde{W}_t \quad (1)$$

where K^Q and Σ are $n \times n$ matrices and $(K\theta)^Q$ is an n -vector. The Q superscript is used to distinguish parameters under the equivalent martingale measure from corresponding parameters under the physical measure. The matrix S_t is diagonal with elements

$$S_{t(ii)} \equiv \sqrt{\alpha_i + \beta_i' X_t} \quad (2)$$

where β_i is an n -vector and α_i is scalar. It is convenient to stack the β_i vectors into the matrix β , where β_i' is row i of β . The scalars α_i are stacked in the n -vector α . The following discussion assumes that the dynamics of (1) are well-defined, which requires that $\alpha_i + \beta_i' X_t$ is nonnegative for all i and all possible X_t . Parameter restrictions that ensure these requirements are in Dai and Singleton (2000).

Denote the time- t price of a zero-coupon bond maturing at time $t + \tau$ as $P(X_t, \tau)$. Duffie and Kan show

$$P(X_t, \tau) = \exp[A(\tau) - B(\tau)X_t] \quad (3)$$

where $A(\tau)$ is a scalar function and $B(\tau)$ is an n -valued function. Thus the bond's yield is affine in the state vector:

$$Y(X_t, \tau) = (1/\tau)[-A(\tau) + B(\tau)'X_t]. \quad (4)$$

The functions $A(\tau)$ and $B(\tau)$ can be calculated numerically by solving a series of ordinary differential equations (ODEs).

1.2. The price of risk and expected returns to bonds

The term structure model is completed by specifying the dynamics of X_t under the physical measure, which is equivalent to specifying the dynamics of the price of risk. Denote the state price deflator by π_t . The relative dynamics of π_t are

$$\frac{d\pi_t}{\pi_t} = -r_t dt - \Lambda_t' dW_t \quad (5)$$

where the vector W_t follows a Brownian motion under the physical measure. Element i of the n -vector Λ_t represents the price of risk associated with the Brownian motion $W_{t,i}$. The dynamics of X_t under the physical measure can be written in terms of Λ_t and the parameters of (1):

$$dX_t = ((K\theta)^Q - K^Q X_t)dt + \Sigma S_t \Lambda_t dt + \Sigma S_t dW_t. \quad (6)$$

Instantaneous bond-price dynamics can be written as

$$\frac{dP(X_t, \tau)}{P(X_t, \tau)} = (r_t + e_{\tau,t})dt + v_{\tau,t}dW_t$$

where $e_{\tau,t}$ denotes the instantaneous expected excess return to holding the bond; the expected return, over r_t , of holding a τ -maturity bond for an instant. An application of Ito's lemma combined with the structure of the ODEs in Duffie and Kan reveals that

$$e_{\tau,t} = -B(\tau)' \Sigma S_t \Lambda_t, \quad (7)$$

$$v_{\tau,t} = -B(\tau)' \Sigma S_t. \quad (8)$$

Equation (7) says that variations over time in expected excess returns are driven by variations in both the volatility matrix S_t and the price of risk vector Λ_t . A fully

parametric model for bond-yield dynamics requires specifying a functional form for Λ_t . This form should be sufficiently flexible to capture the empirically-observed behavior of expected excess returns. Thus to motivate the choice of functional form for Λ_t , we briefly review evidence on the behavior of bond returns.

A large literature documents that expected excess returns to Treasury bonds (over returns to short-term Treasury bills) are, on average, near zero, and vary systematically with the term structure.¹ When the slope of the term structure is steeper than usual, expected excess returns to bonds are high, while expected excess returns are low—often negative—when the slope is less steep. Thus across the maturity spectrum, the ratio of mean expected excess bond returns to the standard deviation of expected excess bond returns is low.

Earlier work has also shown that the shape of the term structure is related to the volatility of yields.² However, the slope—expected return relation is not simply proxying for a volatility—expected return relation. Supporting evidence is in Table I, which reports results from regressions of excess monthly bond returns on the lagged slope of the term structure and lagged yield volatility. Monthly returns to portfolios of Treasury bonds are from the Center for Research in Security Prices. Excess returns to these portfolios are produced by subtracting the contemporaneous return to a three-month Treasury bill. The slope of the term structure is measured by the difference between month-end five-year and three-month zero-coupon yields. The zero-coupon yields are interpolated from coupon bonds using the technique of McCulloch and Kwon (1991), as implemented by Bliss (1997).³ Yield volatility is the standard deviation of the five-year zero-coupon bond’s yield, measured by the the square root of the sum of squared daily changes in the yield during the month.

The sample period is July 1961 through December 1998. The results in Table I reveal that month t ’s volatility has no statistically significant predictive power for excess bond returns in month $t + 1$. By contrast, all of the estimated slope parameters are significant at the ten percent level and half are significant at the five percent level. In addition, the variation in predicted excess returns is large relative to mean excess returns. Consider, for

¹ The literature is too large to cite in full here. Early research includes Fama and Bliss (1987). Two standard references are Fama and French (1989, 1993).

² This literature is also too large to cite in full. In an important paper, Chan, Karolyi, Longstaff, and Sanders (1992) examine the sensitivity of volatility to the level of short-term interest rates. Andersen and Lund (1997) refine their work by decomposing the variation in interest-rate volatility into a component related to the level of short-term interest rates and a stochastic volatility component.

³ I thank Rob Bliss for providing me with the yield data.

example, bonds with maturities between three and four years. The mean excess return is seven basis points per month, while the standard deviation of predictable excess returns is roughly 25 basis points. In results not detailed here, I find that the conclusions are robust to including volatilities of other bond yields as explanatory variables in the regression.

Armed with this information about the empirical behavior of bond returns, we now discuss three alternative parameterizations of Λ_t .

1.3. Completely affine models

Fisher and Gilles (1996) and Dai and Singleton (2000) adopt the following parameterization of Λ_t . Let λ_1 be an n -vector. Then the price of risk vector Λ_t is given by

$$\Lambda_t = S_t \lambda_1. \tag{9}$$

This class nests multifactor versions of Vasicek (1977) and Cox et al. (1985; hereafter CIR). The main reason for the popularity of this form is that the vector $S_t \Lambda_t$ is affine in X_t . This implies affine dynamics for X_t under both the equivalent martingale and physical measures. Affine dynamics of X_t under the physical measure allow for closed-form calculation of various properties of conditional densities of discretely-sampled yields. These properties are discussed in detail in Duffie, Pan, and Singleton (1999) and Singleton (1999). Of less importance is the fact that $\Lambda_t' \Lambda_t$, which is the instantaneous variance of the log state price deflator, is also affine in X_t . This latter property motivates the term “completely affine,” as discussed in the next subsection.

This structure imposes two related limitations on Λ_t . First, variation in the price of risk vector is completely determined by the variation in S_t . Therefore variations in expected excess returns to bonds are driven exclusively by the volatility of yields, an implication that appears inconsistent with the evidence in Table I. Second, the sign of element i of Λ_t is the same as that of element i of λ_1 , because the diagonal elements of S_t are restricted to be nonnegative. The importance of this limitation will be clear in Section 2.

1.4. Essentially affine models

The essentially affine class nests the completely affine class. We first define the elements of a diagonal matrix S_t^- as

$$S_{t(ii)}^- = \begin{cases} (\alpha_i + \beta_i' X_t)^{-1/2}, & \text{if } \inf(\alpha_i + \beta_i' X_t) > 0; \\ 0, & \text{otherwise.} \end{cases}$$

Thus, if diagonal element i of S_t is bounded away from zero, its reciprocal is diagonal element i of S_t^- . For any diagonal element of S_t with a lower bound of zero (whether or not it is accessible given the dynamics of X_t), the associated element of S_t^- is set to zero. Therefore the elements of S_t^- do not explode as the corresponding elements of S_t approach zero.

The form of Λ_t used in the essentially affine model is

$$\Lambda_t = S_t \lambda_1 + S_t^- \lambda_2 X_t \quad (10)$$

where λ_2 is an $n \times n$ matrix. This form shares with (9) two important properties. First, if $S_{t(ii)}$ approaches zero, Λ_t does not go to infinity. Second, $S_t \Lambda_t$ is affine in X_t . Therefore the physical dynamics of X_t are affine, which is convenient for empirical estimation.

There are three important differences between (9) and (10). First, with $\lambda_2 \neq 0$, $\Lambda_t' \Lambda_t$ is not affine in X_t . Therefore this model is not completely affine, but the variance of the state price deflator does not affect bond prices. This is the motivation for the term “essentially affine.” Second, the tight link between the price of risk vector and the volatility matrix is broken. The essentially affine setup allows for independent variation in prices of risk, which is the kind of flexibility needed to fit the empirical behavior of expected excess returns to bonds. Third, the sign restriction on the individual elements of Λ_t is removed.

For future reference we need to explicitly determine the physical dynamics of X_t . Substitute (10) into (6) and define I^- as the $n \times n$ diagonal matrix with $I_{ii}^- = 1$ if $S_{t(ii)}^- \neq 0$, $I_{ii}^- = 0$ if $S_{t(ii)}^- = 0$. Then the physical dynamics in the essentially affine model can be written as

$$dX_t = ((K\theta)^Q - K^Q X_t)dt + \Sigma \left[S_t^2 \lambda_1 + I^- \lambda_2 X_t \right] dt + \Sigma S_t dW_t. \quad (11)$$

Combining terms and denoting element i of λ_1 by λ_{1i} , (11) can be written as

$$dX_t = [K\theta - K X_t]dt + \Sigma S_t dW_t, \quad (12a)$$

where

$$K = K^Q - \Sigma \begin{pmatrix} \lambda_{11} \beta_1' \\ \vdots \\ \lambda_{1n} \beta_n' \end{pmatrix} + \Sigma I^- \lambda_2, \quad (12b)$$

and

$$K\theta = (K\theta)^Q + \Sigma \begin{pmatrix} \alpha_1 \lambda_{11} \\ \vdots \\ \alpha_n \lambda_{1n} \end{pmatrix}. \quad (12c)$$

1.5. An essentially affine example

The following two-factor model illustrates a number of features of the essentially affine model. The instantaneous interest rate r_t follows a Gaussian process and there is some other factor f_t that follows a square-root process. It is convenient to begin by modeling their dynamics under the physical measure. Under this measure, the processes are independent, as in (13):

$$d \begin{pmatrix} f_t \\ r_t \end{pmatrix} = \begin{pmatrix} k_f & 0 \\ 0 & k_r \end{pmatrix} \left(\begin{pmatrix} \bar{f} \\ \bar{r} \end{pmatrix} - \begin{pmatrix} f_t \\ r_t \end{pmatrix} \right) dt + \begin{pmatrix} \sigma_f & 0 \\ 0 & \sigma_r \end{pmatrix} \begin{pmatrix} \sqrt{f_t} & 0 \\ 0 & 1 \end{pmatrix} d \begin{pmatrix} W_{t,1} \\ W_{t,2} \end{pmatrix}. \quad (13)$$

The model is closed with a description of the dynamics of the market price of risk. If we adopt the completely affine version in (9), the result is the classic Vasicek (1977) model for r_t . In such a setup, the variable f_t is irrelevant for bond pricing, and we are left with a standard one-factor Gaussian model.

If, however, we use the essentially affine specification for the market price of risk, the factor f_t can affect bond prices, even though it cannot affect the path of r_t . The reason is that the compensation that investors require to face the risk of $W_{t,2}$ can vary with f_t . The essentially affine model specifies the price of risk Λ_t as

$$\begin{aligned} \Lambda_t &= \begin{pmatrix} \lambda_{11} \sqrt{f_t} \\ \lambda_{12} \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \lambda_{2(11)} & \lambda_{2(12)} \\ \lambda_{2(21)} & \lambda_{2(22)} \end{pmatrix} \begin{pmatrix} f_t \\ r_t \end{pmatrix} \\ &= \begin{pmatrix} \lambda_{11} \sqrt{f_t} \\ \lambda_{12} \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ \lambda_{2(21)} & \lambda_{2(22)} \end{pmatrix} \begin{pmatrix} f_t \\ r_t \end{pmatrix}. \end{aligned}$$

The dynamics of the state price deflator are therefore

$$\frac{d\pi_t}{\pi_t} = -r_t dt - \begin{pmatrix} \lambda_{11} \sqrt{f_t} \\ \lambda_{12} + \lambda_{2(21)} f_t + \lambda_{2(22)} r_t \end{pmatrix}' d \begin{pmatrix} W_{t,1} \\ W_{t,2} \end{pmatrix}.$$

The dynamics of r_t and f_t under the equivalent martingale measure are, from (12a) and (12b),

$$\begin{aligned}
d \begin{pmatrix} f_t \\ r_t \end{pmatrix} &= \begin{pmatrix} k_f + \sigma_f \lambda_{11} & 0 \\ \sigma_r \lambda_{2(21)} & k_r + \sigma_r \lambda_{2(22)} \end{pmatrix} \left(\begin{pmatrix} \bar{f}^Q \\ \bar{r}^Q \end{pmatrix} - \begin{pmatrix} f_t \\ r_t \end{pmatrix} \right) dt + \\
&\quad \begin{pmatrix} \sigma_f & 0 \\ 0 & \sigma_r \end{pmatrix} \begin{pmatrix} \sqrt{f_t} & 0 \\ 0 & 1 \end{pmatrix} d \begin{pmatrix} \tilde{W}_{t,1} \\ \tilde{W}_{t,2} \end{pmatrix}
\end{aligned} \tag{14}$$

where \bar{f}^Q and \bar{r}^Q are the means of f_t and r_t , respectively, under the equivalent martingale measure.

There are three important differences between this description of bond-price dynamics and the standard Vasicek model. First, the current level of the instantaneous interest rate r_t affects the price of interest rate risk, through the parameter $\lambda_{2(22)}$. In Vasicek, the price of interest rate risk is constant. Second, there is a source of uncertainty in bond prices that is independent of the physical dynamics of r_t . The factor f_t affects bond prices through the parameter $\lambda_{2(21)}$. Chacko (1997) builds an affine term structure model expressly designed to exhibit this second feature, and my example was inspired by his (substantially more complicated) model. We will see in Section 4 that this kind of feature is critical to understanding the actual dynamics of Treasury bond yields. Third, the price of risk associated with innovations in $W_{t,2}$ can change sign, depending on the level of the factor f_t .

Because this model takes as a primitive the dynamics of the state-price deflator, it is incapable of providing us a utility-based explanation for sign changes in investors' willingness to face this risk. However, we know from the results of stochastic differential utility that given arbitrary state-price deflator dynamics, there exists some utility gradient and optimal consumption process that are consistent with the deflator dynamics. For a textbook discussion, see Duffie (1996).

The essentially affine structure of Λ_t , although more flexible than the completely affine structure, nonetheless imposes limits on the possible dynamics of bond prices. Note that one element of K , which is the first matrix on the right-hand-side of (13), is the same as the corresponding element of K^Q , which is the first matrix on the right-hand-side of (14). Element (1,2) must be zero under both the physical and equivalent martingale measures. Otherwise, the drift of f_t at $f_t = 0$ could be negative (because it would depend on r_t), which cannot be allowed because $\sqrt{f_t}$ enters into S_t .

To free up this element, and thus allow for a more flexible specification of the price of risk, we could model f_t as a Gaussian process. An example of such a model is Fisher (1998). By contrast, if both f_t and r_t were modeled as square-root diffusion processes, the essentially affine structure of Λ would be identical to the completely affine structure. This illustrates a more general point noted by Duffie and Kan (1996) and Dai and Singleton

(2000), and that we will see in current paper’s empirical work. With affine bond-pricing, there is a tradeoff between constructing a model that can capture complicated dynamics in volatilities and a model that can capture complicated dynamics in expected returns.

1.6. Semi-affine models

Duarte (2000) chooses an alternative generalization of completely affine models. Let λ_0 be an n -vector. The price of risk vector is

$$\Lambda_t = \lambda_0 + S_t \lambda_1.$$

With this form, elements of Λ_t can switch sign over time but they cannot move independently of S_t . As noted in Section 1.4, this latter feature appears inconsistent with the empirical evidence. Thus at first glance it appears that the semi-affine setup allows for some, but not all, of the flexibility of the essentially affine setup. However, there are parameterizations of S_t for which the semi-affine model offers more flexibility than does the essentially affine model. One example is the multifactor CIR model, which is the focus of Duarte’s empirical work. It should be noted that the essentially affine and semi-affine frameworks are nested in the following price of risk vector:

$$\Lambda_t = \lambda_0 + S_t \lambda_1 + S_t^- \lambda_2 X_t.$$

As with the semi-affine form of Λ_t , this more general form implies non-affine dynamics of X_t under the physical measure. Duarte (2000) notes that with non-affine physical dynamics, approximation or simulation techniques are typically necessary to reproduce the properties of discretely-sampled yields.

1.7. A canonical form for essentially affine models

There are a variety of normalizations that can be imposed on affine models. Here I follow the lead of Dai and Singleton’s (2000) canonical completely affine model. They normalize Σ to the identity matrix. They also stack X_t so that, if m of the factors affect the instantaneous variance of X_t (because they enter into one of the diagonal elements of S_t) and $n - m$ of them do not, the m factors in the former group are in positions 1 through m of X_t and the $n - m$ factors in the latter group are in positions $m + 1$ through n . The resulting model is called an $A_m(n)$ model. They also set the first m elements of α to zero and the remaining $n - m$ elements to one. Then their version of (2) is

$$S_{t(ii)} = \begin{cases} \sqrt{X_{t,i}}, & i = 1, \dots, m; \\ \sqrt{1 + \beta'_i X_t}, & i = m + 1, \dots, n \end{cases} \quad (15)$$

where for $i = m + 1, \dots, n$,

$$\beta'_i = (\beta_{i1} \quad \dots \quad \beta_{im} \quad 0 \quad \dots \quad 0).$$

Using their framework, we can write the diagonal elements of S_t^- and I^- as

$$\begin{aligned} S_{t(ii)}^- &= \begin{cases} 0, & i = 1, \dots, m; \\ (1 + \beta'_i X_t)^{-1/2}, & i = m + 1, \dots, n \end{cases} \\ I_{ii}^- &= \begin{cases} 0, & i = 1, \dots, m; \\ 1, & i = m + 1, \dots, n. \end{cases} \end{aligned} \quad (16)$$

Note that in (11) the matrix λ_2 shows up only in the term $I^- \lambda_2 X_t$. Therefore we can normalize the first m rows of λ_2 to zero. Now reconsider (7), the instantaneous expected excess return to holding a bond with remaining maturity τ . From (10), (15), and (16), in the canonical form this can be written as

$$e_{\tau,t} = -B(\tau)' \left[\begin{pmatrix} 0_m \\ \lambda_{1(m+1)} \\ \vdots \\ \lambda_{1n} \end{pmatrix} + \left(\begin{pmatrix} M_{m \times m}^a & 0_{m \times (n-m)} \\ M_{(n-m) \times m}^b & 0_{(n-m) \times (n-m)} \end{pmatrix} + \begin{pmatrix} 0_{m \times n} \\ L_{(n-m) \times n} \end{pmatrix} \right) X_t \right]. \quad (17)$$

In (17), 0_m is an m -vector of zeros. The $0_{p \times q}$ matrices are defined similarly. The submatrix M^a is a diagonal matrix with the i^{th} diagonal element equal to element i of λ_1 . Row i of M^b is given by the first m elements of the vector $\lambda_{1(m+i)} \beta'_{m+i}$. The submatrix L consists of rows $(m + 1)$ through n of λ_2 .

The additional flexibility of the essentially affine model in fitting time-variation in expected excess returns to bonds is captured by the matrix L . In a completely affine setup, L is a zero matrix. Therefore any elements of X_t that do not affect the instantaneous volatility of X_t (i.e., elements $m + 1, \dots, n$), are also incapable of affecting instantaneous expected excess returns to bonds. When L is nonzero, any such elements of X_t can affect expected excess returns. In addition, L provides a mechanism for all other elements in X_t to affect expected returns through a channel other than M^a or M^b .

If all elements of X_t affect the instantaneous volatility (i.e., a correlated multifactor CIR model, or what Dai and Singleton (2000) call an $A_n(N)$ model), there is no L matrix

(it has zero rows). Therefore the essentially affine model generalizes the completely affine model only when there is at least one element in X_t that does not affect the instantaneous volatility of X_t .

2. The intuition behind the failure of completely affine models

A successful model of the term structure should be consistent with the variety of term-structure shapes observed in the data. For example, the model should be capable of producing low, flat term structures; low, steep term structures; and high, inverted term structures. In addition, the model should reproduce the empirically-observed patterns in expected returns to bonds; or equivalently, produce forecasts of future yields that subsume the forecasting information in the slope of the term structure. This section explains that completely affine models fit to the historical behavior of Treasury yields will not simultaneously achieve both of these goals.

For our purposes, the key features of the excess returns to bonds are that they are, on average, small, and exhibit substantial predictable variation. Recall from Section 1 that $e_{\tau,t}$ denotes the instantaneous expected excess return to a bond with maturity τ . Although we do not observe instantaneous returns, the evidence in Table I suggests that the ratio $E(e_{\tau,t})/\sqrt{\text{Var}(e_{\tau,t})}$ is small—well below one—for all τ . (This ratio is the inverse of the coefficient of variation for $e_{\tau,t}$.)

We will see below that completely affine models can be parameterized to produce low values of $E(e_{\tau,t})/\sqrt{\text{Var}(e_{\tau,t})}$ for all τ . However, completely affine models can fit this behavior only by giving up the ability to fit a wide range of term-structure shapes. Conversely, they can be parameterized to fit observed term-structure shapes, but not the behavior of expected excess returns. The intuition underlying this result is best seen in two steps. We will first examine the behavior of one-factor completely affine models. Then we will see that the important properties of one-factor models carry over to multifactor models.

2.1. One-factor models

The intuition in a completely affine one-factor model is straightforward. Expected instantaneous excess bond returns, $e_{\tau,t}$, are proportional to the factor's variance, hence they are bounded by zero. In order for a random variable that is bounded by zero to have a standard deviation substantially larger than its mean, it must be highly skewed. This high skewness is a tight restriction on the admissible values of $e_{\tau,t}$, and thus a tight restriction on the admissible values of the factor.

To see this clearly, we work through the math. Our goal is to reproduce the stylized fact that $E(e_{\tau,t})/\sqrt{\text{Var}(e_{\tau,t})}$ is small. We restrict our attention to a non-Gaussian model, because in a completely affine Gaussian model $\text{Var}(e_{\tau,t}) = 0$. The model is

$$\begin{aligned} r_t &= \delta_0 + x_t, \\ dx_t &= k(\theta - x_t)dt + \sigma\sqrt{x_t}dW_t, \\ \Lambda_t &= \lambda_1\sqrt{x_t}. \end{aligned}$$

From (7), the instantaneous expected excess return to a τ -maturity bond is

$$e_{\tau,t} = -B(\tau)\sigma\lambda_1x_t.$$

Therefore the inverse of the $e_{\tau,t}$'s coefficient of variation is

$$\frac{E(e_{\tau,t})}{\sqrt{\text{Var}(e_{\tau,t})}} = \frac{E(x_t)}{\sqrt{\text{Var}(x_t)}} = \frac{\theta}{\sqrt{\text{Var}(x_t)}}. \quad (18)$$

Equation (18) implicitly imposes $B(\tau)\sigma\lambda_1 < 0$, which is the condition that mean excess bond returns are positive. We set $E(e_{\tau,t})/\sqrt{\text{Var}(e_{\tau,t})} = 0.3$, which is a typical ratio for predictable excess returns in Table I. We set the unconditional mean and standard deviation of the instantaneous interest rate to 5.5 percent and 2.9 percent respectively. These values correspond to the moments of the three-month bill yield over 1952 through 1998. In this model, $\text{Var}(r_t) = \text{Var}(x_t)$. Plugging the standard deviation into (18) produces $\theta = 0.87$ percent. Therefore $\delta_0 = 4.63$ percent to fit the mean instantaneous interest rate.

The requirement that the mean of x_t is small relative to its standard deviation gives the model little flexibility in producing short-term interest rates that are below average. The instantaneous interest rate r_t cannot be less than δ_0 , or 4.63 percent. But over the period 1952 through 1998, three-month yields have ranged from 0.6 to 16 percent. Put differently, the model's parameters and the observed variation in short-term interest rates over this period imply a range of x_t from -4.0 to 11.4 ; the implied x_t is negative in more than 40 percent of the monthly observations. Nonnegativity in implied values of x_t requires $\theta > 4.9$ percent. With such a θ , the model would have the flexibility to fit the wide range of short-term interest rates observed in the data, but the ratio $E(e_{\tau,t})/\sqrt{\text{Var}(e_{\tau,t})}$ would exceed 1.6.

We can also think about this model's restriction on the behavior of interest rates in terms of skewness in expected excess returns. In order to produce a small value of

$E(e_{\tau,t})/\sqrt{\text{Var}(e_{\tau,t})}$, the model will generate expected excess returns that are always positive, usually very close to zero, and occasionally well above zero. But as noted in Section 1, observed expected excess returns are not so positively skewed; they range from positive to negative.

2.2. Multifactor models

Multifactor models are better at fitting the behavior of expected excess bond returns. For example, it is simple to generate a near-zero value of $E(e_{\tau,t})/\sqrt{\text{Var}(e_{\tau,t})}$ for a specified maturity, while retaining substantial flexibility in fitting term structure shapes. In a multifactor CIR model, the only requirement is prices of risk (elements of Λ_t) with different signs. If one element of Λ_t is positive and another negative, then at some maturity the factor loadings will weight these prices of risk such that $E(e_{\tau,t}) = 0$ and $\text{Var}(e_{\tau,t}) > 0$.

However, completely affine models will not produce near-zero values of $E(e_{\tau,t})/\sqrt{\text{Var}(e_{\tau,t})}$ for *all* maturities while also allowing for a wide variety of term structure shapes. To slightly oversimplify, the intuition is that long-maturity bond yields are affected by only one factor—the factor with the greatest persistence under the equivalent martingale measure. Thus we can use the earlier intuition developed for one-factor models to conclude that multifactor models cannot reproduce the observed behavior of long-maturity yields.

The reason why only a single factor will affect long-maturity bond yields is practical, not theoretical. There are a variety of types of shocks that affect the term structure (e.g., level, slope, twist), and multifactor models capture this variety through factors that die out at different rates under the equivalent martingale measure. In principle, we could construct a model with multiple factors affecting long-bond yields. The only requirement is to force the factors to share the same low speed of mean reversion. But by doing so, we weaken the major advantage of multifactor models—the ability to fit different kinds of shocks to the term structure. Thus such a model would produce a poor fit of term structure data relative to a model in which each factor had its own speed of mean reversion.

The failure of completely affine models to fit the empirical behavior of bonds can be seen in the parameter estimates of three-factor completely affine models in Dai and Singleton (2000). They use U.S. dollar interest rate swap yields to estimate the same general three-factor completely affine models that are estimated in this paper. I use the parameters of their preferred model and the swap yields to produce implied time series of the state vector and expected excess returns to bonds. The results of this exercise, which are not reported in any table, indicate the model captures the combination of low mean and high volatility for expected excess returns. However, in over one-quarter of the

observations in their data, the implied value of the state vector violates a nonnegativity constraint. The violations tend to occur when the long end of the term structure is well below its average. Thus the results in Dai and Singleton support the conclusion that completely affine models do not simultaneously fit the behavior of expected excess returns to bonds and the variety of term structure shapes in the data.

3. Estimation of essentially affine models

3.1. Three-factor affine models

All of the affine models estimated in this paper have three underlying factors ($n = 3$). Litterman and Scheinkman (1991) find that three factors explain the vast majority of Treasury bond price movements. This is fortunate, because general three-factor affine models are already computationally difficult to estimate owing to the number of parameters. Adding another factor would have made this investigation impractical. Seven models are estimated: Four completely affine models and three essentially affine models. A completely affine model is estimated for each possible number of factors that do not affect the instantaneous volatility of X_t (from three to zero). The canonical form discussed in Section 1.7 is used. The estimated models are, in their terminology, $A_0(3)$ through $A_3(3)$. The other models that are estimated are the essentially affine generalizations of $A_0(3)$, $A_1(3)$, and $A_2(3)$. (Recall $A_3(3)$ has no essentially affine generalization.)

The estimated models share the following expressions for the instantaneous interest rate, the physical dynamics of X_t , and the price of risk vector:

$$r_t = \delta_0 + \delta_1 X_{t,1} + \delta_2 X_{t,2} + \delta_3 X_{t,3}, \quad (19a)$$

$$d \begin{pmatrix} X_{t,1} \\ X_{t,2} \\ X_{t,3} \end{pmatrix} = \left[\begin{pmatrix} (K\theta)_1 \\ (K\theta)_2 \\ (K\theta)_3 \end{pmatrix} - \begin{pmatrix} k_{11} & k_{12} & k_{13} \\ k_{21} & k_{22} & k_{23} \\ k_{31} & k_{32} & k_{33} \end{pmatrix} \begin{pmatrix} X_{t,1} \\ X_{t,2} \\ X_{t,3} \end{pmatrix} \right] dt + S_t dW_t, \quad (19b)$$

$$S_{t(ii)} = \sqrt{\alpha_i + (\beta_{i1} \ \beta_{i2} \ \beta_{i3}) X_t}, \quad (19c)$$

$$\Lambda_t = S_t \begin{pmatrix} \lambda_{11} \\ \lambda_{12} \\ \lambda_{13} \end{pmatrix} + S_t^- \begin{pmatrix} \lambda_{2(11)} & \lambda_{2(12)} & \lambda_{2(13)} \\ \lambda_{2(21)} & \lambda_{2(22)} & \lambda_{2(23)} \\ \lambda_{2(31)} & \lambda_{2(32)} & \lambda_{2(33)} \end{pmatrix} X_t. \quad (19d)$$

Depending on the model, various restrictions are placed on the parameters in (19a) through (19d).

3.2. *The data*

I use month-end yields on zero-coupon Treasury bonds (interpolated from coupon bonds) calculated using the method in McCulloch and Kwon (1991). Their sample, which ends in February 1991, is extended by Bliss (1997). The entire data set covers the period January 1952 through December 1998.⁴ I restrict my attention to maturities less than or equal to ten years owing to the large number of missing observations for longer-maturity bonds.

In order to perform both in-sample and out-of-sample tests, I estimate term-structure models using data from 1952 through 1994. The final four years of data are reserved for constructing out-of-sample forecast errors.

3.3. *The estimation technique*

I estimate these models using quasi maximum likelihood (QML), which is particularly easy to implement with completely and essentially affine models. Although QML does not use all of the information in the probability density of yields, it fully exploits the information in the first and second conditional moments of the term structure. Thus QML will capture the tension in affine models between fitting conditional means and conditional variances.

Another advantage of QML (which it shares with maximum likelihood and related techniques) is that there is a positive probability that estimated model could actually generate the observed time series of term structures. This is an important concern in estimating affine term structure models. As the discussion in Section 2 highlighted, there is a tradeoff between fitting the coefficients of variation in expected excess bond returns and fitting the variety of term structure shapes observed in the data. A model estimated with QML will guarantee that the time- t state vector implied by time- t yields is in the state vector's admissible space (to avoid a likelihood of zero). By contrast, consider techniques such as Efficient Method of Moments (EMM) that compare sample moments from the data with population moments simulated from the model. These techniques do not require that the estimated term structure model be sufficiently flexible to reproduce the term structure shapes in the data. The parameters of the model in Dai and Singleton (2000), which were estimated with EMM, illustrate this point.

⁴ Bliss and McCulloch-Kwon use slightly different filtering procedures, thus the yields they report over periods of overlapping data do not match exactly. This raises the issue of where to splice the series together. I use the yields in McCulloch and Kwon over their entire sample period and use the Bliss data after February 1991.

I implement QML following Fisher and Gilles (1996), which contains further details. I assume that at each month-end t , $t = 1, \dots, T$, yields on n bonds are measured without error. (Recall n is the dimension of the state vector.) These bonds have fixed times-to-maturity τ_1, \dots, τ_n . Yields on k other bonds are assumed to be measured with serially uncorrelated, mean-zero measurement error. The variance-covariance matrix of this measurement error is the constant matrix V .

To compute the log-likelihood value for a candidate parameter vector, stack the perfectly-observed yields in the vector Y_t and the imperfectly-observed yields in the vector \tilde{Y}_t . Denote the parameter vector by Θ . Given Θ , Y_t can be inverted using (4) to form an implied state vector \hat{X}_t , as in (20).

$$\hat{X}_t = H_1^{-1}(Y_t - H_0). \quad (20)$$

In (20), H_0 is an n -vector with element i given by $A(\tau_i)/\tau_i$ and H_1 is an $n \times n$ matrix with row i given by $B(\tau_i)/\tau_i$. The candidate parameter vector is required to be consistent with Y_t . This is enforced by requiring \hat{X}_t to be in the admissible space for X_t , which is equivalent to requiring that the diagonal elements of S_t in (19c) be real.

Given \hat{X}_t , implied yields for the other k bonds can be calculated. Stack them in \hat{Y}_t . The month t measurement error vector is then $\epsilon_t = \hat{Y}_t - \tilde{Y}_t$. To compute the quasi-likelihood value, assume that the one-period-ahead conditional distribution of the state variables is multivariate normal and equal to

$$f_X(X_{t+1}|X_t).$$

The mean and variance-covariance matrix of X_{t+1} can be computed using the results in the Appendix, thus $f_X(X_{t+1}|X_t)$ is known. Then the distribution of Y_{t+1} conditional on Y_t is

$$f_Y(Y_{t+1}|Y_t) = \frac{1}{|\det(H_1)|} f_X(\hat{X}_{t+1}|\hat{X}_t).$$

Also assume that the measurement error is jointly normally distributed with distribution $f_\epsilon(\epsilon_t)$. The log likelihood of observation t is then

$$l_t(\Theta) = \log f_Y(Y_t|Y_{t-1}) + \log f_\epsilon(\epsilon_t). \quad (21)$$

Stationarity is imposed on the model by requiring that the eigenvalues of K are positive, allowing $f_Y(Y_1|Y_0)$ to be set equal to the unconditional distribution of Y_t . The estimated parameter vector Θ^* is chosen to solve

$$\max_{\Theta} L(\Theta) = \sum_{t=1}^T l_t(\Theta).$$

In the estimation that follows, I assume that the bonds with no measurement error are those with maturities of six months, two years, and ten years. This choice was motivated by the desire to span as much of the term structure as possible without assuming that the three-month yield, which exhibits some idiosyncratic behavior, is observed without error. The bonds with measurement error fill in the gaps in this term structure, with maturities of three months, one year, and five years. The variance-covariance matrix of the measurement error is estimated using its Cholesky decomposition:

$$E(\epsilon_t \epsilon_t') = V = CC'. \quad (22)$$

There are six elements to be estimated in the lower triangular matrix C . An earlier version of the paper assumed a diagonal structure for V . Although the results of the more general structure strongly reject the assumption that these measurement errors are uncorrelated, the parameter estimates of the rest of the model are largely unaffected by the form chosen for V .

It is, however, important to include these additional bond yields in the estimation procedure. Earlier versions of this paper did not include bond yields measured with error. The earlier results showed that the general three-factor models studied here—especially the essentially affine models—can produce wildly implausible term structure shapes. These shapes happen to intersect with observed term structures at three maturities: the maturities associated with the bonds measured without error. By including bonds measured with error, the likelihood values associated with these implausible term structures are penalized heavily.

3.4. *The maximization technique*

The QML functions for these models have a large number of local maxima. The most important reason for this is the lack of structure placed on the feedback matrix K . Similar QML values can be produced by very different interactions among the elements of the state vector. Another reason is that the feasible parameter space is not convex for any model with nonconstant volatilities. A feasible parameter vector satisfies the requirement that the diagonal elements of S_t are real for all t . Because I use the canonical form of Section 1.7, this requirement is satisfied when $\hat{X}_{t,i} \geq 0$ for $i \leq m$. (Recall that m is the number of state variables that affect the instantaneous volatility of X_t .) Therefore

the requirement imposes $m \times T$ restrictions on the parameter vector. The restrictions are nonlinear functions of the parameters and the data. These problems led to the following maximization technique.

Step 1. Randomly generate a parameter vector from a multivariate normal distribution with a diagonal variance-covariance matrix. The means and variances were arbitrarily set to ‘plausible’ values.

Step 2. Use (20) to calculate \hat{X}_t for all t .

Step 3. If the parameter vector is not feasible, return to step 1; otherwise proceed.

Step 4. Use Simplex to determine the parameter vector that maximizes the QML value.

Step 5. Using the final parameter vector from step 4 as a starting point, use NPSOL to make any final improvements in the QML value.

This procedure is repeated until Steps 4 and 5 are reached 1,000 times. For most of the estimated models, there was little improvement in the QML value after the first few hundred iterations.

3.5. Specification tests

These specification tests use the fact that QML estimation can be viewed as a GMM estimator. The moments are the first derivatives of the quasi log likelihood function with respect to the parameter vector, resulting in an exactly-identified model. By imposing overidentifying moment conditions we can test the adequacy of the model.

3.5.1. Tests of nested models

Recall that for $m < n$, the completely affine model $A_m(n)$ is nested in a corresponding essentially affine model. The essentially affine version has an additional $n(n - m)$ free parameters corresponding to the bottom $n - m$ rows in the matrix λ_2 . We can test the null hypothesis that these free parameters are all equal to zero, using the GMM version of a likelihood ratio test. A textbook discussion is in Greene (1997). Define the column vector $h_t(\Theta)$ as the derivative of (21) with respect to the parameter vector Θ , and define the column vector $h(\Theta)$ as the mean of these T vectors. Define the (inverse of) the weighting matrix W_t as the mean of the outer product of the single-period derivatives:

$$W_t^{-1}(\Theta) = (1/T) \sum_{t=1}^T h_t(\Theta)h_t(\Theta)'. \quad (23)$$

Denote the parameter vector for the essentially affine model estimated by QML as Θ_0^* . Denote a restricted parameter vector Θ_1 as a parameter vector that imposes the completely affine restriction on λ_2 . Choose it to solve

$$q = \min_{\Theta_1} Th(\Theta_1)'W_t(\Theta_0^*)h(\Theta_1). \quad (24)$$

The results of Hansen (1982) imply that under the null hypothesis, q is distributed as $\chi^2((n - m)n)$. Similar tests can be used to evaluate other parametric restrictions on the estimated models. These other tests are discussed in more detail in Section 4.

3.5.2. Testing the covariance between forecast errors and the term structure slope

This test asks whether the yield forecasts produced by the estimated models include the information in the slope of the term structure. Given a parameter vector Θ associated with a particular model, the implied state vector $\hat{X}_{t-\Delta}$ is given by inverting yields observed at time $t-\Delta$. The Δ -period-ahead conditional mean $E(X_t|\hat{X}_{t-\Delta})$ can then be constructed. Given this expected state vector, expected Δ -period-ahead bond yields and associated forecast errors can also be constructed. We need some notation for forecast errors. Denote by e_{t,Δ,τ_i} the forecast error realized at time t for a τ_i -maturity bond, where the forecast is made at time $t-\Delta$. The forecast errors for v bonds of different maturities are stacked in the vector $e_{t,\Delta}$.

$$e_{t,\Delta} \equiv (e_{t,\Delta,\tau_1} \quad e_{t,\Delta,\tau_2} \quad \dots \quad e_{t,\Delta,\tau_v})'$$

If an estimated term structure model does not make systematic forecast errors, forecasts of time- t yields made at time $t-\Delta$ should have forecast errors uncorrelated with any variable known at time $t-\Delta$. This motivates the specification test. Denote the slope of the yield curve at time $t-\Delta$ by $s_{t-\Delta}$. If the model is correctly specified,

$$E[(e_{t,\Delta} - \overline{e_{t,\Delta}})(s_{t-\Delta} - \overline{s_{t-\Delta}})] = 0. \quad (25)$$

Equation (25) contains v moment conditions which can be used as overidentifying restrictions in GMM estimation of an affine model. The other moment conditions are standard QML moments, which are the derivatives of (21) with respect to each element of the parameter vector. The weighting matrix is calculated at the QML parameter estimates, which are consistent under the null hypothesis that the model is correctly specified. Then an analogue to q in (24) is calculated. Again from the results of Hansen (1982), this value is distributed as $\chi^2(v)$ under the null hypothesis.

The use of overlapping observations in these moment conditions produces sample moments that exhibit serial correlation. Therefore the weighting matrix does not take the simple form of (23). I experimented with a variant of Hansen and Hodrick (1980)’s weighting matrix, but in practice the matrix was typically not positive definite. I therefore adopted the approach of Newey and West (1987).

To implement this test I set $\Delta = 1/2$, so that six-month-ahead forecasts are examined. This horizon was chosen arbitrarily. A cursory investigation of other forecast horizons indicated that the results of the specification tests were insensitive to this choice. I used eight lags in the Newey-West calculation of the weighting matrix; experimentation with similar lag lengths did not materially affect the results. I set $v = 3$, and form forecasts for maturities of six months, two years, and ten years. (These are the same maturities that are assumed to have no measurement error.) The slope of the term structure is measured by the difference between the five-year bond yield and the three-month bond yield. The first six observations are dropped to account for the length of the forecast horizon.

4. Results

4.1. An overview

Table II reports the QML values for each estimated model. Results for ten models are shown. The first seven model specifications are labeled “unrestricted.” This means that the only parameter restrictions imposed are those implied by the canonical form. These restrictions are either normalizations or requirements of no-arbitrage. To both limit the danger of overfitting and to aid in the interpretation of the parameter estimates, more parsimonious specifications are also estimated. These “preferred” specifications will be discussed after the unrestricted specifications are considered.

Among the unrestricted models, the completely affine $A_0(3)$ model has the lowest QML value, indicating that it provides the worst overall fit to the first and second conditional moments of yields. This is not surprising, given that the model cannot generate time-variation in either expected returns or conditional variances of yields. The completely affine $A_3(3)$ model has the highest QML value, followed closely by the essentially affine and completely affine $A_2(3)$ models. These are the models with the greatest flexibility in fitting variations in conditional volatilities.

The first specification test reveals that the unrestricted completely affine $A_0(3)$ model is overwhelmingly rejected by its more general essentially affine counterpart. The corresponding test statistic for the completely affine $A_1(3)$ is smaller, but still overwhelmingly

rejects the model. The $A_2(3)$ completely affine model has a QML value almost indistinguishable from that of its essentially affine counterpart, and its additional restrictions are not rejected.

The reason for the decrease in the χ^2 statistics as m increases is the decrease in additional flexibility offered by essentially affine models as m increases. Recall from (19d) that the difference between completely and essentially affine models is in the elements of λ_2 . When $m = 0$ all elements of λ_2 are free. When $m = 1$ the top row of λ_2 is set to zero. When $m = 2$ only the bottom row of λ_2 is free. Thus an increase in m provides for greater flexibility in fitting conditional variances of yields but also provides for less flexibility (in an essentially affine model) in fitting expected excess returns to bonds.

Related evidence is reported for the second specification test. Only the essentially affine $A_0(3)$ model passes this test. For all other models, the slope of the term structure is too strongly correlated with the forecast errors to satisfy the null hypothesis that the correlation is zero.

To limit the size of the paper, I report more detailed information for only three of the models. They are the essentially affine $A_0(3)$, $A_1(3)$, and completely affine $A_2(3)$ models. The first is of particular interest because of its forecasting ability, the second illustrates the tradeoff between forecasting ability and fitting conditional variances, while the third is the completely affine model that does the best at forecasting, as measured by the χ^2 statistic on the second specification test.

For each of these models, I estimate specifications that are more parsimonious than the unrestricted specifications. I first computed the t -statistics for the unrestricted parameter estimates. I then set to zero all parameters for which the absolute t -statistics did not exceed one and reestimated the models.⁵ This procedure eliminated five parameters from the completely affine $A_2(3)$ model and seven parameters from both the essentially affine $A_0(3)$ model and the essentially affine $A_1(3)$ model. For each preferred model, a joint test of the parameter restrictions is constructed using an analogue to (24). The test statistics and corresponding p -values are reported in the “First Test Stat” column.⁶

⁵ With the completely affine $A_2(3)$ model, the parameter β_{32} was set to one instead of zero.

⁶ The test statistic for the essentially affine $A_1(3)$ model suggests a rejection of the preferred model in favor of the unrestricted model. However, the large test statistic appears to be a consequence of approximation errors in numerical computation of the derivative of the log-likelihood function with respect to k_{32} . The Numerical Recipes *dfidr* routine (a robust method for calculating derivatives and estimates of errors in the derivatives) reported large errors regardless of the initial stepsize. Because the estimate of parameter in the unrestricted model was nearly zero, and setting it to zero had a negligible effect on the QML likelihood function, I set it to zero in the preferred model.

Parameter estimates for these preferred models are in Tables III through V. To conserve space, parameter estimates for the other models are not reported in the paper, and are available on request. The parameters from (19a) through (19d) are reported along with parameters of the Cholesky decomposition in (22).

Table III reports the parameter estimates for the $A_0(3)$ essentially affine model. The canonical form imposes a lower triangular structure on K and imposes $\alpha = 1$, $\beta = 0$, $K\theta = 0$. Table IV reports the parameter estimates for the $A_1(3)$ essentially affine model. One feature of this table deserves mention. The parameter $(K\theta)_2$ is nonzero, but no standard error is reported. This is the result of two normalizations imposed on the model: $\theta_2 = \theta_3 = 0$. The normalizations are imposed by setting $(K\theta)_2$ and $(K\theta)_3$ to the necessary values given K . Other restrictions imposed in the canonical form are $\alpha_1 = k_{12} = k_{13} = 0$, $\alpha_2 = \alpha_3 = \beta_{11} = 1$, and $\beta_{ij} = 0, i \geq 1, j > 1$. Finally, Table V reports the parameter estimates for the $A_2(3)$ completely affine model. In the canonical form of the $A_2(3)$ model, $\alpha_1 = \alpha_2 = \beta_{33} = 0$, $\alpha_3 = \beta_{11} = \beta_{22} = 1$, $\beta_{ij} = 0$ for $i < 3, i \neq j$, and $\lambda_2 = 0$. The preferred specification sets $\beta_{31} = 0$ and $\beta_{32} = 1$, so that the second state variable drives the conditional volatilities of both the second and third state variables. The element $(K\theta)_3$ is nonzero with no standard error because $\theta_3 = 0$ in the canonical model.

4.2. An analysis of forecast errors

The estimated models, combined with month t bond yields, can be used to construct forecasts of month $t+i$ bond yields. Here we examine the accuracy of these forecasts, both in-sample and out-of-sample. The in-sample period is January 1952 through December 1994. The out-of-sample period is January 1995 through December 1998. We focus on bonds with maturities of six months, two years, and ten years, and forecast horizons of three, six, and twelve months. Forecast accuracy is measured by the root mean squared forecast error (RMSE). In-sample RMSEs are reported in Table VI and out-of-sample RMSEs are reported in Table VIII. In Tables VII (in-sample) and IX (out-of-sample), forecast errors are regressed on the slope of the yield curve to determine whether the forecast method captures the forecasting power of the slope.

We need benchmarks to use in evaluating forecast accuracy. The simplest benchmark is a random walk. The month t yield on a τ -maturity bond is used as a forecast of the month $t+i$ yield on a τ -maturity bond. The RMSEs associated with this forecast method are reported in the ‘‘RW’’ columns of Tables VI and VIII. Note that the tables report different patterns in RMSEs across bonds. In the earlier period, yields were more volatile, with volatility declining with maturity. In the later period, yield volatility was higher at

long maturities than at short maturities. Thus the out-of-sample period should provide a good test of the robustness of the estimated affine models.

A more sophisticated benchmark uses OLS regressions that predict future changes in yields with the current slope of the term structure. The regression is

$$Y_{\tau,t+i} - Y_{\tau,t} = b_0 + b_1(Y_{5yr,t} - Y_{3mo,t}) + e_{\tau,t+i}. \quad (26)$$

The parameters of (26) are estimated using in-sample data. The equation is then used to construct forecasts and forecast errors for both the in-sample and out-of-sample periods. The resulting RMSEs are in the columns labeled “OLS” in Table VI and Table VIII. Although the in-sample RMSE for the regression is guaranteed to be no larger than the random walk RMSE, that is not true out-of-sample. Indeed, for eight of the nine combinations of maturity and forecast horizon, the out-of-sample OLS RMSE exceeds that of the random walk.

The in-sample parameter estimates from (26) are reported in Table VII in the column labeled “RW.” This may seem like a misprint (why aren’t they labeled “OLS”?), but recall that Table VII reports the parameter estimates of regressions of forecast errors on the month t slope of the yield curve. With the random walk method of forecasting, the regression examined in Table VII is identical to the regression used to produce OLS forecasts. The results document that short-maturity yields tend to rise and long-maturity yields tend to fall when the slope is steeper than average, although the statistical evidence at the short end is weak. These results correspond to the standard violations of the expectations hypothesis of interest rates.

This violation is also apparent in the behavior of bond yields in the out-of-sample period. The “RW” column in Table IX reports the results of estimating (26) from January 1995 through December 1998. The point estimates are typically more negative than their counterparts in Table VII, although the t -statistics are smaller owing to fewer observations.

The final six columns in Table VI through Table IX examine the forecasting ability of various affine models. The results document that the completely affine $A_2(3)$ model is a failure at forecasting future interest rates. Table VI reports that in-sample, both the unrestricted and preferred specifications produce forecasts that are worse than those produced by the assumption that yields follow random walks. This unimpressive performance is mirrored by the performance of the other completely affine models examined in this paper. For every estimated model, the assumption that yields follow a random walk produces superior in-sample forecasts for each of these maturities and forecast horizons. (These additional results are not reported in any table.)

The regressions reported in Table VII show that the forecast errors of the completely affine $A_2(3)$ model are strongly negatively correlated with the slope of the term structure. The parameter estimates are more negative than are the corresponding parameter estimates in the random walk case. The model completely misses the forecasting information in the slope of the term structure. When the term structure is more steeply sloped than usual, the OLS forecast is that long-maturity yields will fall, but the model forecasts that the yields will rise. Put differently, the model is consistent with the expectations hypothesis, and the observed bond yields are not.

This poor forecasting performance carries over to the out-of-sample period. Table VIII documents that the unrestricted specification produces forecasts that are inferior to random-walk forecasts in five of the nine combinations of maturity and forecast horizon. The preferred specification does even worse, producing inferior forecasts for seven of the nine combinations. The point estimates in Table IX confirm that the model's forecasts get the wrong sign of the relationship between the slope of the term structure and future changes in yields.

The essentially affine models produce dramatically better forecasts. The most successful forecasting model, both in-sample and out-of-sample, is the essentially affine, completely Gaussian model. Table VI documents that within the sample, both the unrestricted and preferred $A_0(3)$ models outforecast the OLS regressions (and therefore also outforecast the random walk assumption) for each combination of maturity and forecast horizon. Table VIII makes the same point out-of-sample. Moreover, these forecasts capture the predictive power of the term-structure slope. In Tables VII and IX, the only evidence for predictability of forecast errors is in the out-of-sample forecast errors for ten-year bonds at the twelve month horizon.

The essentially affine $A_1(3)$ model is not quite as successful as the Gaussian model at forecasting. From Table VI, we see that in-sample forecasts from both the unrestricted and preferred specifications are typically superior to random-walk forecasts, but outforecast OLS regressions for only half of the maturity/horizon combinations. Moreover, from Table VIII, the forecast errors are negatively correlated with the slope of the yield curve. The statistical strength of this negative correlation rises as both the bond's maturity and the forecasting horizon lengthen.

An examination of Table VIII indicates that this essentially affine model performs somewhat better out-of-sample. Forecasts from the preferred specification are superior to random-walk and OLS forecasts at all maturities and forecast horizons. Nonetheless, Table IX indicates that the model's out-of-sample forecast errors are negatively correlated with the slope of the yield curve. Thus the model misses some of the explanatory power of the

term-structure slope.

The general pattern in these forecasts is that forecasting ability falls as the ability to fit time-variation in conditional variances rises. The $A_0(3)$ model, although its forecasting ability is superior, has the lowest QML value among the three considered here because it implies that yield volatilities are constant through time. Notwithstanding its poor forecasting performance, the $A_2(3)$ model has a much higher QML value than either of the essentially affine models discussed here because of its ability to fit variations in volatilities. In the next subsection we consider more explicitly the tradeoff between forecasting power and time-varying volatilities.

4.3. The predictability of excess returns and volatilities

A few diagrams help shed light on the behavior of these competing models. Figure 1 is a graphical summary of the behavior of the preferred essentially affine $A_0(3)$ model. Panel A displays instantaneous effects that one-standard-deviation shocks to each factor have on the term structure of yields. The three shocks can be interpreted as a level shock (the long dashes), a slope shock (the solid line), and a twist (the short dashes). Panel B displays the (nonexistent) instantaneous effect of these shocks on yield variances.

Panel C displays the effect that these shocks have on bonds' instantaneous expected excess returns (over r_t). There are two distinct types of shocks to expected returns. The short dashes correspond to the twist shock in Panel A. This shock has a strong effect on instantaneous expected returns, but it is also very short-lived. (This latter fact cannot be seen in the panel.) Thus this shock is responsible for high-frequency fluctuations in expected excess returns.

The other type of shock to expected excess returns corresponds to the slope shock in Panel A. It is more persistent (this also cannot be seen in the panel), and thus accounts for more persistent fluctuations in expected returns. The combined effects of these shocks on expected excess returns to two-year bonds are displayed in Panel E. Panel F is the same plot for ten-year bonds. These latter panels show that expected excess returns fluctuate sharply and widely around zero. For example, the expected instantaneous excess return in Panel E has a mean of 1.25 percent and a standard deviation of 3.09 percent.

Because this model is so successful at forecasting future yields, it is worth a more careful examination. An intuitive way to interpret shocks to bond yields is to decompose the shocks into shocks to expected future short-term interest rates and shocks to expected excess returns. This decomposition is straightforward, thus I will not discuss it in detail here. Instead, I will simply summarize the results.

A positive level shock corresponds to an immediate, near-permanent increase in short-term interest rates. The half-life of the shock to short-term interest rates is more than eleven years. Because the shock does not substantially alter investors' required excess returns to bonds, short-maturity and long-maturity bond yields respond in the same way to this shock.

A positive slope shock corresponds to an immediate increase in short-term interest rates that lasts about as long as a business cycle. The half-life of the shock is four years. Because short-term interest rates are expected to decline over time, the shock lowers the slope of the term structure. The shock also lowers expected excess returns to bonds by affecting the price of risk vector. We can see this in the parameters of λ_2 in Table III. An increase in the first factor (the slope factor) affects the price of risk of the third factor (the level factor) through element (3,1) of λ_2 . This decrease in expected returns further decreases the slope of the term structure because longer-maturity bond returns are more sensitive than shorter-maturity bond returns to level shocks, and thus to the price of risk of level shocks.

Twists are very similar to the “ f_t ” factor in the two-factor example discussed in Section 2. A twist shock has basically no effect on current or future short-term interest rates. Instead, the shock changes investors' required excess returns to bonds by affecting the price of risk associated with the level and slope factors. The half-life of such a shock is less than three months. We can call this a “flight to quality” shock. Investors experience short-lived periods of unwillingness to hold risky Treasury instruments, thus driving expected excess bond returns higher.

Figure 2 contains information about the preferred essentially affine $A_1(3)$ model. Panel A displays a level shock, a slope shock, and a twist shock. The solid line is the level shock, and it affects the conditional variance of yields, as shown in Panel B. The long-dashed line is the twist shock, and in Panel C we see its strong effect on expected excess returns. However, Panel C also indicates that the other two shocks have little effect on expected excess returns. The net effect is that in Panels E and F, the fluctuations in expected excess returns are less volatile than the fluctuations in the corresponding panels in Figure 1. For example, the expected instantaneous excess return in Panel E has a mean of 1.90 percent and a standard deviation of 1.85 percent.

Why does a shock to the slope affect expected excess returns in Figure 1 but not in Figure 2? The answer is that the channel that operates in the model underlying Figure 1 is unavailable in the model underlying Figure 2. Panel C in Figure 1 reflects a relationship between shocks to the slope and shocks to the price of risk of level shocks. These cross-factor relationships are more limited in the essentially affine $A_1(3)$ model. In the canonical

form, the first factor drives conditional volatilities, thus its price of risk cannot be affected by any other factors. Figure 2 indicates that this first factor is the level factor; shocks to the slope cannot affect its price of risk. Therefore this model produces poorer forecasts of future bond yields than does the essentially affine $A_0(3)$ model.

Figure 3 displays the same panels for the preferred completely affine $A_2(3)$ model. The model generates a richer pattern of time-variation in volatilities than do the other two models. The cost of these more accurate measures of volatility is an inability to fit expected excess returns. Expected excess returns in Panels E and F are always positive, never large, and not volatile. For example, the expected instantaneous excess return in Panel E has a mean of 0.79 percent and a standard deviation of 0.41 percent. Moreover, these expected excess returns roughly track the instantaneous interest rate displayed in Panel D. Because higher short-term rates typically correspond to lower slopes, the figure indicates that expected excess returns move inversely with the slope of the yield curve; but this is counterfactual.

The results discussed in this section indicate that the completely affine $A_2(3)$ model fails to reproduce the behavior of expected excess returns to Treasury bonds. The same conclusion holds for the other completely affine models estimated in this paper that are not discussed in detail here. The models systematically fail to capture the large fluctuations in expected excess returns to bonds. Essentially affine models do a better job of reproducing the behavior of expected excess returns, although the magnitude of the improvement is inversely related to the ability of the models to fit the time-variation in conditional variances of yields.

5. Concluding comments

Recent term structure research has concentrated on what I call completely affine models. This paper documents that completely affine models do not forecast future yields well over the nearly fifty year period examined here. They consistently underestimate future returns to bonds when the term structure is more steeply-sloped than usual; put differently, these models do not reproduce the well-known failure of the expectations hypothesis.

Essentially affine models generalize completely affine models. They allow greater flexibility in fitting variations in the price of interest rate risk over time, while retaining the affine time-series and cross-sectional properties of bond prices. One of the essentially affine models investigated in this paper—the pure Gaussian model—generates reasonable forecasts of future yields, in the sense that the predictive power of the term structure is subsumed within the model’s forecasts.

The forecast accuracy of this Gaussian model allows us to properly interpret the usual level, slope, and twist yield-curve factors in terms of their predictions for future short-term interest rates and excess returns to longer-term bonds. Level shocks correspond to near-permanent changes in interest rates and only minimal changes in expected excess returns. Slope shocks correspond to business-cycle-length fluctuations in both interest rates and expected excess returns to bonds, while twist shocks correspond to short-lived “flight to quality” variations in expected excess returns. In other words, twist shocks do not affect current or expected future short-term interest rates; they are pure shocks to risk premia.

Essentially affine models are not magic bullets. The models cannot capture time-variation in conditional variances without giving up part of their flexibility in fitting time-variation in the price of interest rate risk. It remains to be seen whether an essentially affine model can be constructed that reproduces the time-variation observed in both the conditional variances of yields and expected returns to bonds.

Appendix. First and second moments of affine state variables

This appendix describes closed-form representations for first and second conditional moments of a state vector that follow the affine process of (12a) and (2). The results are an application (and a specialization) of the results in Fisher and Gilles (1996).

Assume that K can be diagonalized, or

$$K = NDN^{-1}, \quad D \text{ diagonal.} \quad (A1)$$

The diagonal elements of D are denoted d_1, \dots, d_n . A discussion of computing moments when K cannot be diagonalized is in Fisher and Gilles (1996).

The approach taken here is to compute the first and second conditional moments of a linear transformation of X_t . The transformation is chosen so that the feedback matrix K is diagonal under the transformation. The linear transformation is then reversed to calculate the conditional moments of X_t . Define

$$X_t^* \equiv N^{-1}X_t. \quad (A2)$$

Then the dynamics of X_t^* are, from (12a), (2), (A1), and (A2),

$$dX_t^* = D(\theta^* - X_t^*) + \Sigma^* S_t^* dW_t, \quad (A3)$$

where

$$\begin{aligned} S_{t(i,i)}^* &= \sqrt{\alpha_i + \beta_i^{*'} X_t^*}, \\ \theta^* &= N^{-1}\theta, \\ \Sigma^* &= N^{-1}\Sigma, \\ \beta^* &= \beta N. \end{aligned}$$

We now calculate the first and second moments of X_t^* . Some notation is helpful. If Z is an n -vector, the $n \times n$ diagonal matrix in which element (i, i) equals Z_i is denoted $diag(Z)$. If Z is a diagonal matrix, the diagonal matrix in which element (i, i) equals $e^{Z_{ii}}$ is denoted e^Z . Finally, the n -vector $\beta_{\bullet i}$ is column i of β .

1. Conditional mean

The expectation of X_T^* conditional on X_t^* is given by

$$E[X_T^*|X_t^*] = \theta^* + e^{-D(T-t)}(X_t^* - \theta^*). \quad (A4)$$

Because $e^{-D(T-t)}$ is diagonal, this expectation can also be simply expressed element-by-element:

$$E[X_T^*|X_t^*] = \theta_i^* + e^{-d_i(T-t)}(X_{t,i}^* - \theta_i^*). \quad (A4')$$

Another useful way to express (A4) is by separating the terms that depend on X_t^* from the terms that do not:

$$E[X_T^*|X_t^*] = \left(I - e^{-D(T-t)}\right) \theta^* + e^{-D(T-t)} X_t^*. \quad (A4'')$$

Given this conditional mean of X_T^* , we reverse the transformation to express the conditional mean of X_T .

$$E[X_T|X_t] = NE[X_T^*|X_t] = N \left(I - e^{-D(T-t)}\right) \theta^* + Ne^{-D(T-t)} N^{-1} X_t.$$

Note that the conditional mean of X_T could be expressed directly in terms of the parameters of (12a); no transformation into X_t^* is required, because the above expression is equivalent to

$$E[X_T|X_t] = \left(I - e^{-K(T-t)}\right) \theta + e^{-K(T-t)} X_t$$

where $e^{-K(T-t)}$ is the fundamental matrix associated with $-K(T-t)$. The value of the approach taken here is that (A4') is used in determining the conditional variance-covariance matrix of X_t .

2. Conditional variance

The matrix $\Sigma^* S_t^* S_t^{*'} \Sigma^{*'}$ is the instantaneous variance-covariance matrix of the transformed state vector. We can write this as

$$\begin{aligned} \Sigma^* S_t^* S_t^{*'} \Sigma^{*'} &= \Sigma^* \text{diag}(\alpha^*) \Sigma^{*'} + \sum_{i=1}^n \Sigma^* \text{diag}(\beta_{\bullet i}^*) \Sigma^{*'} X_{t,i}^* \\ &\equiv G_0 + \sum_{i=1}^n G_i X_{t,i}^* \end{aligned} \quad (A5)$$

where $G_0 \equiv \Sigma^* \text{diag}(\alpha^*) \Sigma^{*'}$ and the $n \times n$ matrices G_i are defined as $[\Sigma^* \text{diag}(\beta_{\bullet i}^*) \Sigma^{*'}]$. Define the $n \times n$ matrix $F(t, s)$ as

$$F(t, s) \equiv G_0 + \sum_{i=1}^n G_i [E(X_s^* | X_t^*)]_i.$$

This matrix is the instantaneous variance-covariance matrix of X_s^* , but evaluated at the expectation of X_s^* (conditional on time- t information) instead of at the true value of X_s^* . Using (A4'), this matrix can be expressed as

$$F(t, s) = G_0 + \sum_{i=1}^n G_i [\theta_i^* + e^{-d_i(s-t)}(X_{t,i}^* - \theta_i^*)]. \quad (A6)$$

Fisher and Gilles show the conditional variance of X_T^* can be written as

$$\text{Var}[X_T^* | X_t^*] = \int_t^T e^{-D(T-s)} F(t, s) e^{-D(T-s)} ds. \quad (A7)$$

Substituting (A6) into (A7) produces (A8):

$$\begin{aligned} \text{Var}[X_T^* | X_t^*] &= \int_t^T e^{-D(T-s)} G_0 e^{-D(T-s)} ds + \\ &\sum_{i=1}^n \left[\theta_i^* \int_t^T e^{-D(T-s)} G_i e^{-D(T-s)} ds \right] + \\ &\sum_{i=1}^n \left[(X_{t,i}^* - \theta_i^*) \int_t^T e^{-D(T-s)} G_i e^{-D(T-s)} e^{-d_i(s-t)} ds \right]. \end{aligned} \quad (A8)$$

If $f(j, k)$ maps (j, k) into the scalar value f , the notation $\{f(j, k)\}$ denotes the matrix with element (j, k) given by $f(j, k)$. The conditional variance can then be written as

$$\begin{aligned} \text{Var}[X_T^* | X_t^*] &= \int_t^T \left\{ [G_0]_{j,k} e^{(s-T)(d_j+d_k)} \right\} ds + \\ &\sum_{i=1}^n \left[\theta_i^* \int_t^T \left\{ [G_i]_{j,k} e^{(s-T)(d_j+d_k)} \right\} ds \right] + \\ &\sum_{i=1}^n \left[(X_{t,i}^* - \theta_i^*) \int_t^T \left\{ [G_i]_{j,k} e^{(s-T)(d_j+d_k) - d_i(s-t)} \right\} ds \right]. \end{aligned} \quad (A9)$$

Integrating (A9) produces (A10).

$$\begin{aligned}
Var[X_T^*|X_t^*] &= \left\{ (d_j + d_k)^{-1} [G_0]_{j,k} \left(1 - e^{-(T-t)(d_j+d_k)} \right) \right\} \\
&\quad \sum_{i=1}^n \left[\theta_i^* \left\{ (d_j + d_k)^{-1} [G_i]_{j,k} \left(1 - e^{-(T-t)(d_j+d_k)} \right) \right\} \right] + \\
&\quad \sum_{i=1}^n \left[(X_{t,i}^* - \theta_i^*) \left\{ (d_j + d_k - d_i)^{-1} [G_i]_{j,k} (e^{-d_i(T-t)} - e^{-(d_j+d_k)(T-t)}) \right\} \right]
\end{aligned} \tag{A10}$$

Note that by collecting terms, the variance-covariance matrix in (A10) can be rewritten in terms of the individual elements of X_t^* as in (A11)

$$Var[X_T^*|X_t^*] = b_0 + \sum_{i=1}^n b_i X_{t,i}^*. \tag{A11}$$

The $n \times n$ matrices $b_i, i = 0, \dots, n$ depend on the horizon $T - t$. We now calculate the conditional variance of X_T using the notation of (A11). Since

$$Var(X_T|\Omega) = NVar(X_T^*|\Omega)N'$$

we have

$$Var(X_T|X_t) = Nb_0N' + \sum_{i=1}^n \left(\sum_{j=1}^n Nb_jN'N_{j,i}^{-1} \right) X_{t,i}.$$

References

- Andersen, Torben G., and Jesper Lund, 1997, "Estimating continuous-time stochastic volatility models of the short-term interest rate," *Journal of Econometrics* 77, 343-377.
- Bliss, Robert R., 1997, "Testing term structure estimation methods," *Advances in Futures and Options Research* 9, 197-231.
- Campbell, John Y., and Robert J. Shiller, 1991, "Yield spreads and interest rate movements: A bird's eye view," *Review of Economic Studies* 58, 495-514.
- Chacko, George, 1997, "Multifactor interest rate dynamics and their implications for bond pricing," Working paper, Harvard Business School (Boston, MA).
- Chan, K.C., G. Andrew Karolyi, Francis A. Longstaff, and Anthony B. Sanders, 1992, "An empirical comparison of alternative models of the short-term interest rate," *Journal of Finance* 47, 1209-1227.
- Cox, John C., Jonathan E. Ingersoll, and Stephen A. Ross, 1985, "A theory of the term structure of interest rates," *Econometrica* 53, 385-407.
- Dai, Qiang, and Kenneth J. Singleton, 2000, "Specification analysis of affine term structure models," *Journal of Finance* 55, 1943-1978.
- Duarte, Jefferson, 2000, "The relevance of the price of risk in affine term structure models," Working paper, University of Chicago (Chicago, IL).
- Duffie, Darrell, 1996, *Dynamic Asset Pricing Theory*, 2nd Ed., Princeton University Press (Princeton, NJ).
- Duffie, Darrell, and Rui Kan, 1996, "A yield-factor model of interest rates," *Mathematical Finance*, 379-406.
- Duffie, Darrell, Jun Pan, and Kenneth Singleton, 1999, "Transform analysis and asset pricing for affine jump-diffusions," Working paper, Stanford University (Stanford, CA).
- Fama, Eugene F., and Robert R. Bliss, 1987, "The information in long-maturity forward rates," *American Economic Review* 77, 680-692.

- Fama, Eugene F., and Kenneth R. French, 1989, "Business conditions and expected returns on stocks and bonds," *Journal of Financial Economics* 25, 23-49.
- Fama, Eugene F., and Kenneth R. French, 1993, "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics* 33, 3-56.
- Fisher, Mark, 1998, "A simple model of the failure of the expectations hypothesis," Working paper, Federal Reserve Board (Washington, DC).
- Fisher, Mark, and Christian Gilles, 1996, "Estimating exponential-affine models of the term structure," Working paper, Federal Reserve Board (Washington, DC).
- Greene, William H., 1997, *Econometric Analysis*, 3rd Ed., Prentice-Hall (Upper Saddle River, NJ).
- Hansen, Lars Peter, 1982, "Large sample properties of generalized method of moment estimators," *Econometrica* 50, 1029-1053.
- Hansen, Lars Peter, and Robert J. Hodrick, 1980, "Forward exchange rates as optimal predictors of future spot rates: An econometric analysis," *Journal of Political Economy* 88, 829-853.
- Litterman, Robert and Jose Scheinkman, 1991, "Common factors affecting bond returns," *Journal of Fixed Income* 1, 54-61.
- McCulloch, J. Huston, and Heon-Chul Kwon, 1993, "U.S. Term Structure Data, 1947-1991," Ohio State University working paper 93-6 (Columbus, OH).
- Newey, Whitney K., and Kenneth D. West, 1987, "A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix," *Econometrica* 55, 703-708.
- Singleton, Kenneth J., 1999, "Estimation of affine pricing models using the empirical characteristic function," Working paper, Stanford University (Stanford, CA).
- Vasicek, Oldrich, 1977, "An equilibrium characterization of the term structure," *Journal of Financial Economics* 5, 177-188.

Table I
Regressions of excess returns to Treasury bonds
July 1961 through December 1998

Monthly excess returns to portfolios of Treasury coupon bonds are regressed on the previous month's term-structure slope and an estimate of the interest rate volatility during the previous month. The slope of the term structure is measured by the difference between five-year and three-month zero-coupon yields (interpolated from coupon bonds). Monthly volatility is measured by the square root of the sum of squared daily changes in the five-year zero-coupon bond yield. Asymptotic t-statistics, adjusted for generalized heteroskedasticity, are in parentheses. There are 449 monthly observations.

Maturity (years)	Mean excess return (%)	_____Coef on _____ slope volatility		Std. dev. of fitted excess rets
$0 < m \leq 1$	0.011	0.027 (1.76)	0.116 (0.96)	0.036
$1 < m \leq 2$	0.045	0.085 (1.85)	0.413 (1.27)	0.119
$2 < m \leq 3$	0.064	0.132 (1.88)	0.582 (1.20)	0.179
$3 < m \leq 4$	0.074	0.187 (2.38)	0.706 (1.35)	0.241
$4 < m \leq 5$	0.063	0.214 (2.37)	0.692 (1.16)	0.265
$5 < m \leq 10$	0.094	0.296 (2.69)	0.804 (1.08)	0.354

Table II
Statistical comparison of estimated models

Three-factor affine models are estimated with quasi maximum likelihood (QML). The data consist of month-end yields on zero-coupon bonds with maturities between three months and ten years, from January 1952 through December 1994. The models differ in the number of factors m that affect the instantaneous variance of yields and in the flexibility of the price of risk parameterization. Essentially affine models allow the price of risk to vary independently from the instantaneous variance of yields, while completely affine models do not. “Unrestricted” models impose no restrictions on the parameters other than those required by no-arbitrage. “Preferred” models drop parameters that contribute little to the models’ QML values.

Two specification tests are reported. The first is of the null hypothesis that the model’s parameter restrictions are true. For “unrestricted” models, the test compares completely affine models to their more general essentially affine counterparts. For “preferred” models, the test compares the preferred model to its unrestricted counterpart. The second tests the null hypothesis that the six-month-ahead yield forecast errors for bonds of three different maturities are uncorrelated with the slope of the term structure at the time the forecasts are made. Under the null, the test statistics are distributed as $\chi^2(\text{number of param restrictions})$ and $\chi^2(3)$ respectively.

Model Type	m	Number of Free Params	QML value	First Test Stat (p-value)	Second Test Stat (p-value)
Unrestricted					
Completely	0	19	15171.94	62.689 (0.000)	12.297 (0.006)
Completely	1	23	15380.31	26.133 (0.000)	18.521 (0.000)
Completely	2	24	15395.74	0.860 (0.835)	9.938 (0.019)
Completely	3	25	15396.34		33.482 (0.000)
Essentially	0	28	15196.45		2.385 (0.596)
Essentially	1	29	15392.47		16.639 (0.001)
Essentially	2	27	15396.04		11.381 (0.010)
Preferred					
Completely	2	19	15393.55	1.238 (0.941)	10.406 (0.015)
Essentially	0	21	15190.68	3.443 (0.841)	1.449 (0.694)
Essentially	1	22	15387.91	17.882 (0.013)	14.361 (0.002)

Table III
Parameter estimates for the preferred essentially affine $A_0(3)$ model

The model is defined in equation (19). With this version of the model, α is a vector of ones and both β and $K\theta$ are identically zero. The matrix C is the Cholesky decomposition $V = CC'$ of the variance-covariance matrix of the cross-sectional errors in fitting yields on bonds with maturities of three months, one year, and five years. Parameters are estimated with QML. Asymptotic standard errors are in parentheses.

Parameter	Constant term		
	1	2	3
	δ_0 0.044 (0.025)		
	Index number (i)		
	1	2	3
δ_i	0.01895 (0.00223)	0.00790 (0.00218)	0.00992 (0.00051)
k_{1i}	0.564 (0.047)	0	0
k_{2i}	0	3.257 (0.672)	0
k_{3i}	-0.545 (0.202)	0	0.062 (0.051)
λ_{1i}	-0.625 (0.146)	-0.235 (0.099)	-0.207 (0.057)
$\lambda_{2(1i)}$	0	1.742 (0.254)	0
$\lambda_{2(2i)}$	0	-1.711 (0.717)	0
$\lambda_{2(3i)}$	0.648 (0.206)	0.297 (0.186)	-0.061 (0.051)
C_{1i}	0.00227 (0.00013)	0	0
C_{2i}	-0.00050 (0.00007)	0.00084 (0.00004)	0
C_{3i}	0	-0.00017 (0.00006)	0.00093 (0.00004)

Table IV
Parameter estimates for the preferred essentially affine $A_1(3)$ model

The model is defined in equation (19). With this version of the model, $\alpha_1 = \beta_{12} = \beta_{13} = 0$, $\alpha_2 = \alpha_3 = \beta_{11} = 1$, and the first row of λ_2 is zero. The matrix C is the Cholesky decomposition $V = CC'$ of the variance-covariance matrix of the cross-sectional errors in fitting yields on bonds with maturities of three months, one year, and five years. Parameters are estimated with QML. Asymptotic standard errors are in parentheses.

Parameter	Constant term		
	1	2	3
	δ_0	0.014 (0.005)	
	Index number (i)		
	1	2	3
δ_i	0.00088 (0.00021)	0.00118 (0.00053)	0.00256 (0.00124)
$(K\theta)_i$	0.155 (0.048)	-1.910	0
k_{1i}	0.031 (0.020)	0	0
k_{2i}	-0.383 (0.235)	0.594 (0.053)	5.340 (3.833)
k_{3i}	0	0	2.832 (0.490)
β_{2i}	10.269 (9.96)	0	0
β_{3i}	0.291 (0.261)	0	0
λ_{1i}	-0.042 (0.020)	-3.844 (2.415)	0
$\lambda_{2(2i)}$	39.334 (53.816)	0	5.259 (3.647)
$\lambda_{2(3i)}$	0	0	-1.311 (0.565)
C_{1i}	0.00227 (0.00013)	0	0
C_{2i}	-0.00049 (0.00007)	0.00084 (0.00004)	0
C_{3i}	0	-0.00016 (0.00006)	0.00094 (0.00004)

Table V
Parameter estimates for the preferred completely affine $A_2(3)$ model

The model is defined in equation (19). With this version of the model, $\alpha_1 = \alpha_2 = 0$, $\alpha_3 = \beta_{11} = \beta_{22} = \beta_{32} = 1$, the remaining elements of the β matrix are zero, and λ_2 is a matrix of zeros. The matrix C is the Cholesky decomposition $V = CC'$ of the variance-covariance matrix of the cross-sectional errors in fitting yields on bonds with maturities of three months, one year, and five years. Parameters are estimated with QML. Asymptotic standard errors are in parentheses.

Parameter	Constant term		
	1	2	3
δ_0		0.018 (0.004)	
	Index number	(i)	
δ_i	0.00066 (0.00021)	0.00136 (0.00050)	0.00598 (0.00077)
$(K\theta)_i$	0	0.222 (0.103)	-2.299
k_{1i}	0.172 (0.064)	-0.295 (0.056)	0
k_{2i}	-0.197 (0.066)	0.406 (0.059)	0
k_{3i}	0.564 (0.279)	-1.669 (0.234)	1.721 (0.176)
λ_{1i}	-0.042 (0.018)	0	-0.208 (0.058)
C_{1i}	0.00227 (0.00013)	0	0
C_{2i}	-0.00049 (0.00007)	0.00084 (0.00004)	0
C_{3i}	0	-0.00017 (0.00006)	0.00094 (0.00004)

Table VI
Comparison of in-sample forecasting performance

This table reports root mean squared errors (RMSE) for month t forecasts of month $t + i$ bond yields. Eight different forecast methods are compared. The column labeled “RW” (random walk) uses month t yields as forecasts of future yields. The column labeled “OLS” uses a univariate OLS regression to form forecasts, where the dependent variable is change in the yield from t to $t + i$ and the regressor is the month t slope of the yield curve. The final six columns use either completely affine (C.A.) or essentially affine (E.A.) three-factor models to form forecasts. Preferred models are restricted versions of unrestricted models. The models differ in the number of factors j that are allowed to affect conditional volatility ($A_j(3)$).

The regression and affine models are estimated using data from January 1952 through December 1994 and the forecasts are produced over the same period (in-sample forecasts). The slope of the yield curve is the five-year zero-coupon yield less the three-month zero-coupon yield. Bond yields are measured in decimal form (i.e., 0.04 corresponds to 4 percent/year).

Bond Maturity	Forecast Horizon	RW	OLS	Unrestricted			Preferred		
				C. A. $A_2(3)$	E. A. $A_0(3)$	E. A. $A_1(3)$	C. A. $A_2(3)$	E. A. $A_0(3)$	E. A. $A_1(3)$
6 mo	3	1.023	1.020	1.045	1.009	1.019	1.048	1.009	1.019
2 yr	3	0.871	0.869	0.880	0.837	0.847	0.883	0.837	0.853
10 yr	3	0.549	0.532	0.554	0.526	0.543	0.554	0.528	0.547
6 mo	6	1.376	1.370	1.418	1.342	1.367	1.427	1.345	1.368
2 yr	6	1.154	1.149	1.173	1.091	1.121	1.181	1.089	1.133
10 yr	6	0.760	0.722	0.774	0.711	0.756	0.772	0.713	0.764
6 mo	12	1.803	1.797	1.843	1.731	1.798	1.868	1.742	1.798
2 yr	12	1.541	1.529	1.566	1.450	1.527	1.583	1.445	1.544
10 yr	12	1.109	1.018	1.137	1.011	1.121	1.131	1.009	1.133

Table VII
The relation between in-sample forecast errors and the yield-curve slope

Various models are used to produce month t forecasts of month $t + i$ bond yields and the corresponding forecast errors are constructed. This table reports parameter estimates from regressions of forecast errors on the month t slope of the yield curve. Six forecast methods are compared. The column labeled “RW” (random walk) uses month t yields as forecasts of future yields. Therefore the forecast error regression is simply a regression of changes in bond yields from t to $t + i$ on the month t yield-curve slope. The final six columns use either completely affine (C.A.) or essentially affine (E.A.) three-factor models to form forecasts. Preferred models are restricted versions of unrestricted models. The models differ in the number of factors j that are allowed to affect conditional volatility ($A_j(3)$).

The regression and affine models are estimated using data from January 1952 through December 1994 and the forecasts are produced over the same period (in-sample forecasts). The slope of the yield curve is the five-year zero-coupon yield less the three-month zero-coupon yield. Asymptotic t-statistics, in parentheses, are adjusted for generalized heteroskedasticity and moving average residuals.

Bond Maturity	Forecast Horizon	RW	Unrestricted			Preferred		
			C. A. $A_2(3)$	E. A. $A_0(3)$	$A_1(3)$	C. A. $A_2(3)$	E.A. $A_0(3)$	$A_1(3)$
6 mo	3	0.072 (0.73)	-0.182 (-1.84)	-0.041 (-0.42)	-0.135 (-1.39)	-0.182 (-1.83)	0.019 (0.19)	-0.124 (-1.27)
2 yr	3	-0.043 (-0.52)	-0.182 (-2.22)	-0.043 (-0.54)	-0.134 (-1.69)	-0.183 (-2.22)	0.013 (0.16)	-0.129 (-1.63)
10 yr	3	-0.125 (-2.72)	-0.159 (-3.50)	-0.027 (-0.61)	-0.141 (-3.19)	-0.158 (-3.49)	-0.018 (-0.39)	-0.140 (-3.15)
6 mo	6	0.118 (0.91)	-0.324 (-2.55)	-0.085 (-0.69)	-0.252 (-2.03)	-0.326 (-2.53)	0.015 (0.12)	-0.233 (-1.88)
2 yr	6	-0.082 (-0.76)	-0.326 (-3.17)	-0.091 (-0.91)	-0.261 (-2.60)	-0.330 (-3.17)	-0.003 (-0.03)	-0.249 (-2.51)
10 yr	6	-0.220 (-3.45)	-0.280 (-4.48)	-0.049 (-0.78)	-0.255 (-4.14)	-0.280 (-4.46)	-0.031 (-0.50)	-0.252 (-4.09)
6 mo	12	0.129 (0.70)	-0.567 (-3.30)	-0.208 (-1.21)	-0.484 (-2.74)	-0.575 (-3.30)	-0.058 (-0.33)	-0.453 (-2.60)
2 yr	12	-0.158 (-1.06)	-0.551 (-3.86)	-0.191 (-1.32)	-0.486 (-3.26)	-0.560 (-3.86)	-0.069 (-0.48)	-0.462 (-3.15)
10 yr	12	-0.410 (-3.62)	-0.506 (-4.53)	-0.135 (-1.22)	-0.480 (-4.24)	-0.507 (-4.52)	-0.101 (-0.92)	-0.472 (-4.19)

Table VIII
Comparison of out-of-sample forecasting performance

This table reports root mean squared errors (RMSE) for month t forecasts of month $t + i$ bond yields. Eight different forecast methods are compared. The column labeled “RW” (random walk) uses month t yields as forecasts of future yields. The column labeled “OLS” uses a univariate OLS regression to form forecasts, where the dependent variable is change in the yield from t to $t + i$ and the regressor is the month t slope of the yield curve. The final six columns use either completely affine (C.A.) or essentially affine (E.A.) three-factor models to form forecasts. Preferred models are restricted versions of unrestricted models. The models differ in the number of factors j that are allowed to affect conditional volatility ($A_j(3)$).

The regression and affine models are estimated using data from January 1952 through December 1994, while the forecasts are produced over January 1995 through December 1998 (out-of-sample forecasts). For each bond there are $48 - i$ forecasts and associated errors. The slope of the yield curve is the five-year zero-coupon yield less the three-month zero-coupon yield. Bond yields are measured in decimal form (i.e., 0.04 corresponds to 4 percent/year).

Bond Maturity	Forecast Horizon	RW	OLS	Unrestricted			Preferred		
				C. A. $A_2(3)$	E. A. $A_0(3)$	E. A. $A_1(3)$	C. A. $A_2(3)$	E. A. $A_0(3)$	E. A. $A_1(3)$
6 mo	3	0.298	0.298	0.325	0.281	0.288	0.350	0.281	0.284
2 yr	3	0.499	0.511	0.501	0.454	0.458	0.523	0.457	0.450
10 yr	3	0.484	0.498	0.476	0.460	0.457	0.485	0.469	0.453
6 mo	6	0.400	0.413	0.483	0.373	0.399	0.548	0.365	0.385
2 yr	6	0.652	0.675	0.656	0.565	0.576	0.711	0.566	0.560
10 yr	6	0.669	0.693	0.647	0.623	0.616	0.669	0.636	0.606
6 mo	12	0.484	0.523	0.621	0.434	0.488	0.778	0.421	0.455
2 yr	12	0.762	0.787	0.759	0.608	0.635	0.879	0.600	0.606
10 yr	12	0.815	0.829	0.764	0.724	0.719	0.811	0.738	0.698

Table IX

The relation between out-of-sample forecast errors and the yield-curve slope

Various models are used to produce month t forecasts of month $t + i$ bond yields and the corresponding forecast errors are constructed. This table reports parameter estimates from regressions of forecast errors on the month t slope of the yield curve. Six forecast methods are compared. The column labeled “RW” (random walk) uses month t yields as forecasts of future yields. Therefore the forecast error regression is simply a regression of changes in bond yields from t to $t + i$ on the month t yield-curve slope. The final six columns use either completely affine (C.A.) or essentially affine (E.A.) three-factor models to form forecasts. Preferred models are restricted versions of unrestricted models; model parameters that add little to the model’s QML value are set to zero. The models differ in the number of factors j that are allowed to affect conditional volatility ($A_j(3)$).

The regression and affine models are estimated using data from January 1952 through December 1994, while the forecasts are produced over January 1995 through December 1998 (out-of-sample forecasts). For each bond, there are $48 - i$ forecasts and associated errors. The slope of the yield curve is the five-year zero-coupon yield less the three-month zero-coupon yield. Asymptotic t -statistics, in parentheses, are adjusted for generalized heteroskedasticity and moving average residuals.

Bond Maturity	Forecast Horizon	RW	Unrestricted			Preferred		
			C. A. $A_2(3)$	E. A. $A_0(3)$	$A_1(3)$	C. A. $A_2(3)$	E.A. $A_0(3)$	$A_1(3)$
6 mo	3	0.121 (1.04)	-0.206 (-1.71)	0.039 (0.35)	-0.074 (-0.66)	-0.220 (-1.82)	0.113 (1.01)	-0.064 (-0.57)
2 yr	3	-0.151 (-0.76)	-0.268 (-1.37)	-0.002 (-0.01)	-0.108 (-0.59)	-0.284 (-1.45)	0.056 (0.30)	-0.104 (-0.57)
10 yr	3	-0.265 (-1.42)	-0.280 (-1.51)	-0.107 (-0.59)	-0.220 (-1.23)	-0.286 (-1.55)	-0.112 (-0.61)	-0.219 (-1.22)
6 mo	6	0.034 (0.20)	-0.497 (-2.48)	-0.094 (-0.56)	-0.272 (-1.65)	-0.525 (-2.59)	0.014 (0.09)	-0.258 (-1.56)
2 yr	6	-0.380 (-1.11)	-0.560 (-1.68)	-0.153 (-0.52)	-0.324 (-1.10)	-0.588 (-1.75)	-0.071 (-0.24)	-0.318 (-1.08)
10 yr	6	-0.552 (-1.60)	-0.571 (-1.68)	-0.305 (-0.92)	-0.485 (-1.49)	-0.583 (-1.71)	-0.307 (-0.91)	-0.483 (-1.49)
6 mo	12	-0.086 (-0.35)	-0.825 (-3.22)	-0.245 (-1.09)	-0.482 (-2.28)	-0.882 (-3.39)	-0.116 (-0.52)	-0.468 (-2.19)
2 yr	12	-0.844 (-2.27)	-1.037 (-3.00)	-0.500 (-1.55)	-0.734 (-2.39)	-1.088 (-3.13)	-0.405 (-1.24)	-0.727 (-2.35)
10 yr	12	-1.085 (-3.75)	-1.083 (-3.89)	-0.737 (-2.58)	-0.977 (-3.68)	-1.105 (-3.95)	-0.730 (-2.52)	-0.975 (-3.66)




Figure 1. Summary of the estimated essentially affine $A_0(3)$ model

Panels A through C display the instantaneous responses of yields, variances, and expected excess returns (over r_t) to one-standard-deviation shocks to each of the three factors. Panels D through F display fitted expected instantaneous returns over the sample period January 1952 through December 1994. Panel D is the instantaneous interest rate. Panels E and F are the instantaneous expected excess returns to the two-year and ten-year bonds.




Figure 2. Summary of the estimated essentially affine $A_1(3)$ model

Panels A through C display the instantaneous responses of yields, variances, and expected excess returns (over r_t) to one-standard-deviation shocks to each of the three factors. Panels D through F display fitted expected instantaneous returns over the sample period January 1952 through December 1994. Panel D is the instantaneous interest rate. Panels E and F are the instantaneous expected excess returns to the two-year and ten-year bonds.

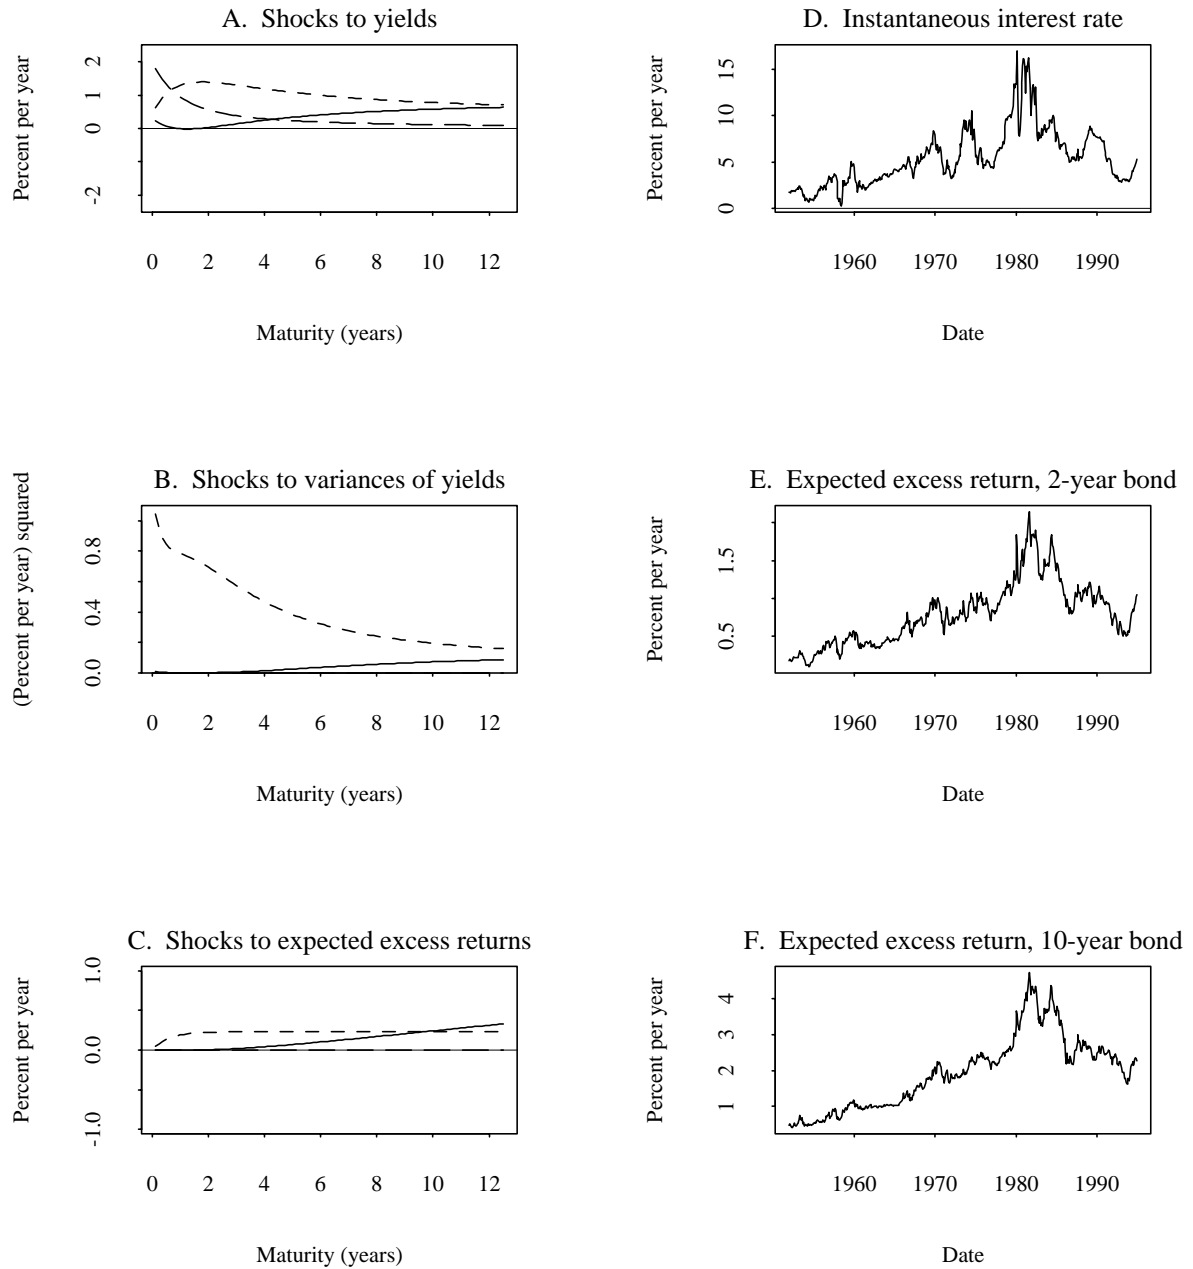


Figure 3. Summary of the estimated completely affine $A_2(3)$ model

Panels A through C display the instantaneous responses of yields, variances, and expected excess returns (over r_t) to one-standard-deviation shocks to each of the three factors. Panels D through F display fitted expected instantaneous returns over the sample period January 1952 through December 1994. Panel D is the instantaneous interest rate. Panels E and F are the instantaneous expected excess returns to the two-year and ten-year bonds.